

WHITE PAPER

5 signs you've outgrown Couchbase

Snapshot

Why replace Couchbase?



It's getting too expensive to scale



You need proven performance for mixed workloads



You want broader multi-site support



You want to use both graph and vector for your AI/ML projects



You need predictable clustering operations when things change

Introduction

If you're like many Couchbase users, you probably found the firm's database a good initial choice for storing and querying JSON data. Its memory-first approach offers some performance benefits, and its Couchbase Lite and Sync Gateway options provide interesting mobile capabilities. However, initial enthusiasm for Couchbase can fade as data volumes and transaction throughputs grow, creating operational difficulties, higher-than-expected ownership costs, and performance issues that can lead to service-level agreement (SLA) violations.

What are 5 signs that your organization may have outgrown Couchbase?

1

It's getting too expensive to scale

- Has Couchbase's memory-first architecture priced you out of some use cases?
- Are you faced with sprawling server footprints?
- Will Couchbase price you out of building new revenue-generating applications?

2

You need proven performance at scale for mixed workloads

- Are Couchbase latency spikes making linear scalability impossible?
- Is write amplification inhibiting your performance goals?
- Does unpredictable performance during busy times threaten your brand reputation?

3

You want broader multi-site support

- Can you benefit from having both synchronous and asynchronous active-active replication options?
- Would you like to define a single cluster that spans multiple geographies or cloud zones with strong consistency?
- Do you want your applications to expand with your business regardless of geography?

4

You want to use both graph and vector data models for you AI/ML projects

- Has Couchbase's document-centric vision limited its usefulness?
- Can you benefit from a single system that is optimized to deliver real-time graph and vector data for AI and graph RAG use cases?
- Could you consolidate your database platforms to better achieve your AI/ML goals?

5

You need predictable clustering operations when things change

- Is your application resilience compromised during cluster rebalancing?
- Is your staff spending too much time and effort managing increasingly large clusters to meet your performance needs?
- Would you like your database choices to satisfy a broader range of business requirements?

1. It's getting too expensive to scale

As data volumes grow and application needs evolve, it's increasingly critical to have a cost-efficient, real-time database that scales easily. Yet, all too often, Couchbase customers find that their infrastructures fall short. Couchbase's architecture can be too costly to deploy for certain use cases, and scaling Couchbase clusters can lead to sprawling footprints that drive up operational costs.

Couchbase was created by merging codebases from Membase (created by members of the original Memcache team) and Apache CouchDB. Couchbase is a NoSQL distributed document database that features a memory-first architecture to achieve high performance, automatically managing a caching layer that keeps frequently accessed data in memory. Memory is allocated on a per-node basis, and different nodes can be configured to run different services (e.g., analytics, text search, data, indexing, query, eventing, and backup).

This approach is effective to a certain extent, but managing the "working set" of data in memory—ideally at least 20%—while storing the rest on disk can result in large clusters with numerous nodes. This is especially true when using the default storage engine to handle data volumes ranging from a few terabytes to petabytes. As data volumes grow and workloads become more varied, cache misses become more likely, leading to unpredictable data access latencies. Cost overruns and service level agreement (SLA) violations often follow.

By contrast, Aerospike employs a unified, efficient storage engine format. It enables firms to choose from in-memory, all-flash, Hybrid Memory Architecture (DRAM and SSDs/flash), and networked NVMe-compatible block storage options, such as cloud-native solutions like AWS EBS or enterprise NAS systems like NetApp. Because each storage category has different price-performance characteristics, Aerospike's approach allows firms to choose the configuration that best suits their needs. Different Aerospike clusters can be configured with different storage options while providing a consistent and familiar experience for developers and operators. In addition, [recent Aerospike optimizations](#) provide even greater efficiency for cloud-native storage (such as EBS) and mixed workloads that have slow or low write operations. A wide variety of publicly available [benchmarks](#) from Aerospike and its customers attest to the system's speed and cost efficiency.

Aerospike deployments typically require 80% fewer nodes and computing resources. This translates to lower TCO, less maintenance, and reduced operational complexity.

2. You need proven performance at scale for mixed workloads

Achieving fast, predictable performance for mixed workloads is essential for databases that support real-time applications. Architectural differences between Couchbase and Aerospike explain why latency spikes and write amplification issues are more likely with Couchbase.

Couchbase's memory-first architecture means that cache misses are inevitable as data volumes grow and evolving application requirements cause data access patterns to vary. Cache misses contribute to latency spikes, particularly at p99 or greater levels. The less data there is in the cache, the higher the cache hit probability. So if the cache miss is 60%, then from roughly the p60 percentile, they will see latencies spike up.

Consider the experience of one firm that benchmarked Couchbase and Aerospike for a workload consisting of 180k writes/second and 75k reads/second over 25 billion objects. As the figure below shows, latency spikes emerged with Couchbase at p99 and became much more pronounced at p999. By contrast, Aerospike's latencies were much lower and more consistent, even at p999.




Figure 1: Customer comparison of upsert latency with Couchbase (left) and Aerospike (right).

What's behind such variation? Differences in storage technologies are a driving factor.

Couchbase can be configured with one of two storage engines: Couchstore (in-memory data store, the default) and Magma (disk storage positioned for larger persistent data sets). Couchstore uses a B-tree-based structure, and certain aspects of this engine can introduce write overhead: e.g., block compression isn't supported, and compaction is single-threaded and not incremental. To address these issues and provide better support for large datasets that don't fit into memory, Couchbase introduced its Magma engine in 2022. Magma combines LSM trees and a segment log approach from log-structured file systems. As of this writing, little performance data from customers in production is available for Magma.

By contrast, Aerospike provides an efficient, unified storage model that spans in-memory, all-flash, and hybrid configurations (indexes in memory and user data on flash). This approach promotes a consistent development experience and delivers RAM-like performance with flash storage for cost efficiency and performance predictability. Aerospike employs a custom file format that bypasses file system, block, and page caching layers, treating flash storage as raw devices. The result is fast, predictable performance at scale, as evidenced by many [customer testimonials](#) and publicly available benchmarks. While Magma is intended to enable Couchbase to serve very large datasets on disk, it does not feature storage driver optimizations comparable to Aerospike's.

3. You want broader multi-site support

Global business operations demand an IT infrastructure that supports fast, predictable, and reliable access to real-time data from multiple geographies. While both Couchbase and Aerospike offer data replication for varied business needs (such as continuous operations, fast localized data access, and disaster recovery), Aerospike offers a broader range of deployment options for building geographically distributed applications.

Aerospike supports both synchronous and asynchronous active-active data replication, and it can be configured across different sites, data centers, availability zones, regions, and cloud providers simultaneously. This allows firms to readily fulfill the resiliency, availability, and accuracy requirements of a broad set of applications.

Couchbase supports asynchronous active-active and active-passive data replication across multiple sites through its [Cross Data Center Replication \(XDRC\)](#) offering. Aerospike provides similar support for asynchronous active-active and active-passive replication via [Cross Datacenter Replication \(XDR\)](#). With Aerospike XDR, replication across clusters is often achieved in sub-millisecond or single-digit milliseconds. XDR also supports selective replication (i.e., data filtering) and performance optimizations to minimize the transfer of frequently updated data. In addition to XDR, Aerospike supports synchronous active-active replication with [multi-site clustering \(MSC\)](#), whereas Couchbase only offers synchronous active-passive replication. Aerospike MSC is accomplished via rack awareness, pegging primary and replica partitions to distinct data centers and racks. With MSC, a single Aerospike cluster can span multiple geographies, and strong, immediate data consistency is automatically enforced. Such capability is particularly useful for global transaction processing.

4. You want the freedom to use graph and vector data models for AI

As customer needs and competitive landscapes change, a flexible data management infrastructure is essential. Support for modern data models such as graph and vector is becoming increasingly important.

Couchbase users model their data as JSON-based documents, each of which can have varied schemas. Users can model key-value data as JSON documents if desired. Couchbase also supports vector search via this model, but not graph data. Increasingly, both graph and vector data models will be used in tandem in [graph RAG](#) architectures.

By contrast, Aerospike supports graph and vector data models, as well as key-value, document, and geospatial models. Data is stored in sets of records, each of which contains a key and one or more named fields (“bins”). Each bin can hold simple or complex data, including collection data types (CDTs) that support fast processing of JSON data. Aerospike’s considerable schema flexibility and efficient support of varied data models enables the database to be used for a broader range of application needs, simplifying the data management infrastructure and reducing the range of required staff skills.

5. You need predictable clustering operations when things change

Simplifying operations and controlling operational overhead is essential as workloads and data volumes increase or fluxuate. That's why self-managing features, modest server footprints, and support for a wide range of real-time applications are critical features of a modern data platform.

Both Couchbase and Aerospike can be deployed as clusters, and both automatically detect changes in cluster status. However, Couchbase clusters require manual rebalancing when nodes are added or removed (unless using Kubernetes). Manual rebalancing increases the risk of reduced resiliency until a rebalance is triggered when the cluster is running with just one copy in cases when the replication factor is set at two (RF = 2). In this case, copies are only created when a rebalance is triggered. Until that period, resiliency is compromised.

By contrast, Aerospike automatically redistributes data when a change in cluster status is detected, reducing staff effort while ensuring strong runtime performance and high data availability. Aerospike's efficient use of computing resources simplifies operations because the number of nodes needed to support a given workload is often substantially less than what Couchbase requires. Fewer nodes mean less effort to monitor and manage the environment and fewer hardware failures.

Consolidating data management technologies is another way firms can simplify operations and cut costs. As mentioned, Aerospike supports a wider range of data models and storage options than Couchbase, making it easier for firms to standardize on Aerospike for a broader range of applications. Many firms initially deploy Aerospike as a cache to promote real-time access to other systems of record or systems of engagement and later leverage Aerospike's built-in persistence features to support additional applications.

Get ready for the future

A real-time, multi-model data management system, Aerospike delivers ultra-fast, predictable performance for read/write workloads at terabyte to petabyte scale with little overhead or complexity. High availability, exceptional scalability, and resource efficiency are other hallmarks of Aerospike's technology. Firms in finance, technology, telecommunications, retail, and other industries use Aerospike for systems of engagement and systems of record, simplifying their infrastructures and often [cutting server footprints by 50-80%](#).

Compared to Couchbase, Aerospike is far less reliant on RAM for lightning-fast performance due to HMA, efficient exploitation of SSDs and other modern hardware technologies, and a wider range of storage management options. Aerospike also features automatic rebalancing of clusters when a change is detected, support for more varied data models, and multi-site clustering technology that enables a single cluster to span multiple geographies or cloud zones.

If you're struggling to achieve what you want with Couchbase or have experienced any of the five signs just discussed, why not explore what Aerospike can do for you? [Contact Aerospike](#) to estimate TCO savings for your workload, or [try Aerospike](#) and see how you can benefit, too.

Playtika



Online gaming

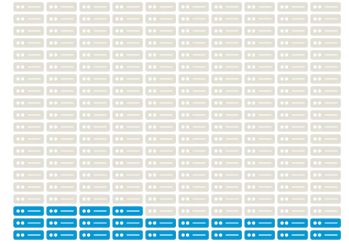
Couchbase challenge: Growth in user demands strained their data management infrastructure, leading to performance problems, unbalanced utilization across clusters, sprawling server footprints, and rampant maintenance issues.

[Read more](#)

200



24



300%

Boost in performance

85%

Server reduction

\$4.2M

Projected savings in TCO over 3 years

Criteo




Real-time ad bidding


Couchbase challenge: Restricted to an outdated version of Couchbase due to scaling and budgetary issues that came with other versions, Criteo struggled to keep up with growing data demands.

[Read more](#)

3200



800



75%

Server reduction

99th

Percentile reads on 290M TPS


AppsFlyer

Mobile marketing analytics and attribution platform

AppsFlyer's traffic trends are very dynamic and unpredictable. Their infrastructure must scale up or down very rapidly. For example, a sports or shopping campaign can increase traffic by 100 percent in one day.

Couchbase challenge: Poor stability led to performance issues and frequent failures, with backups taking over a day to complete.

[Read more](#)



300%

Boost in performance

85%

Server reduction

\$4.2M

Projected savings in TCO
over 3 years

About Aerospike

Aerospike is the real-time database built for infinite scale, speed, and savings. Our customers are ready for what's next with the lowest latency and the highest throughput data platform. Cloud and AI-forward, we empower leading organizations like Adobe, Airtel, Criteo, DBS Bank, Experian, PayPal, Snap, and Sony Interactive Entertainment. Headquartered in Mountain View, California, our offices include London, Bangalore, and Tel Aviv.

For more information, please visit <https://www.aerospike.com>.